

CLAIMS

1 1. A method for avoiding livelock among two or more input/output (I/O) devices
2 of a symmetrical multiprocessor computer system comprising a plurality of intercon-
3 nected processors, one or more shared memories coupled to the processors, and at least
4 one I/O bridge in communicating relationship with the two or more I/O devices, the proc-
5 essors and the one or more shared memories, the method comprising the steps of:

6 providing at least one coherent buffer and at least one non-coherent buffer at the
7 I/O bridge, the non-coherent buffer coupled to the at least one coherent buffer and to at
8 least one of the I/O devices;

9 receiving a request from a first I/O device coupled to the I/O bridge for informa-
10 tion;

11 storing the device requested information in the coherent buffer of the I/O bridge;

12 receiving a system message at the I/O bridge requesting the information stored in
13 the coherent buffer, the system message originating from other than the first I/O device;

14 copying at least a portion of the stored information to the non-coherent buffer;

15 invalidating the stored information within the coherent buffer; and

16 supplying to the first I/O device at least some of the stored information copied
17 into the non-coherent buffer.

1 2. The method of claim 1 wherein

2 the first I/O device is coupled to the non-coherent buffer by an I/O bus having a
3 bus cycle specifying a predetermined number of bits per I/O bus cycle, and

4 the stored information supplied to the first I/O device from the non-coherent
5 buffer is the predetermined number of bits of one bus cycle.

1 3. The method of claim 2 further comprising the steps of:

2 receiving a second request at the I/O bridge from the first I/O device requesting
3 information;

4 determining whether the information of the second request is stored in the coher-
5 ent buffer; and

6 if the information of the second request is stored in the coherent buffer, supplying
7 at least some of the information to the first I/O device.

1 4. The method of claim 3 further comprising the steps of:

2 if the information of the second request is not stored in the coherent buffer, de-
3 termining whether the information of the second request is stored in the non-coherent
4 buffer; and

5 if the information of the second request is stored in the non-coherent buffer, sup-
6 plying the predetermined number of bits of one bus cycle of the information to the first
7 I/O device.

1 5. The method of claim 4 further comprising the steps of:

2 granting the I/O bridge exclusive ownership relative to the plurality of processors
3 and the other I/O bridges of the computer system over the information stored by the I/O
4 bridge; and

5 following the step of invalidating, generating an acknowledgement confirming
6 that the stored information has been invalidated by the I/O bridge.

1 6. The method of claim 5 further comprising the steps of:

2 organizing information stored in the one or more shared memories of the com-
3 puter system into respective cache lines; and

4 providing one or more cache coherency directories, the one or more cache coher-
5 ency directories configured to store an ownership status for each cache line,

6 wherein

7 the system message requesting information originates from one or more of the
8 directories and the acknowledgement is sent to one or more of the directories.

1 7. An input/output (I/O) bridge for use in a distributed shared memory computer

2 system comprising a plurality of interconnected processors and one or more shared

3 memories that are coupled to the processors, the I/O bridge configured to provide inter-

4 communication between one or more I/O devices and the plurality of processors or
5 shared memories, the I/O bridge comprising:
6 at least one coherent buffer configured to store information requested by a first
7 I/O device coupled to the I/O bridge;
8 at least one non-coherent buffer coupled to the coherent buffer and to the one or
9 more I/O devices; and
10 a controller coupled to the coherent buffer and the non-coherent buffer, the con-
11 troller configured to:
12 store at least a portion of the information stored in the coherent buffer in
13 the non-coherent buffer in response to receiving a system message originating
14 from other than the first I/O device requesting the information stored in the coher-
15 ent buffer,
16 invalidate the information within the coherent buffer, and
17 supply to the first I/O device at least some of the information copied into
18 the non-coherent buffer.

1 8. The I/O bridge of claim 7 further wherein
2 the first I/O device is coupled to the non-coherent buffer by an I/O bus having a
3 bus cycle specifying a predetermined number of bits per I/O bus cycle, and
4 the information supplied to the first I/O device from the non-coherent buffer is the
5 predetermined number of bits of one bus cycle.